Constructing the ASCI Computational Grid

Distributed Resource Management (DRM) is developing and deploying software products to support the ASCI mission. The goal is to simplify how users interact with high performance computing (HPC) resources. As part of the computational infrastructure, DRM software will aid in the discovery, reservation, allocation, monitoring, and control of geographically distributed computational resources throughout the nuclear weapons complex. These resources consist of heterogeneous compute nodes, communication, storage, visualization, data, and software. DRM will help realize resource allocation and usage policies.

DRM Goals

DRM will provide the resource management infrastructure for the ASCI grid. The infrastructure will allow for the policy-based management of scarce resources and the hiding of complexity. Such management will maximize throughput and utilization, allow the routine sharing and aggregation of resources, and provide access to additional (currently unmanaged) resources such as visualization capabilities and network quality of service. The infrastructure will be highly resource- and application-independent and will allow the increase in capability and the addition of emerging technologies.

DRM Grid Services

DRM is built on the Globus Resource Management Toolkit. The toolkit was chosen because of its basic functionality, widespread use, and extensibility. We worked closely with the Globus development team to deploy Globus in our isolated testbed (see Figure 1). Using Globus allowed us to install the Nimrod-G scalable parameter study tool in the testbed. Globus also provides utilities for monitoring grid state which is shown in Figure 2.

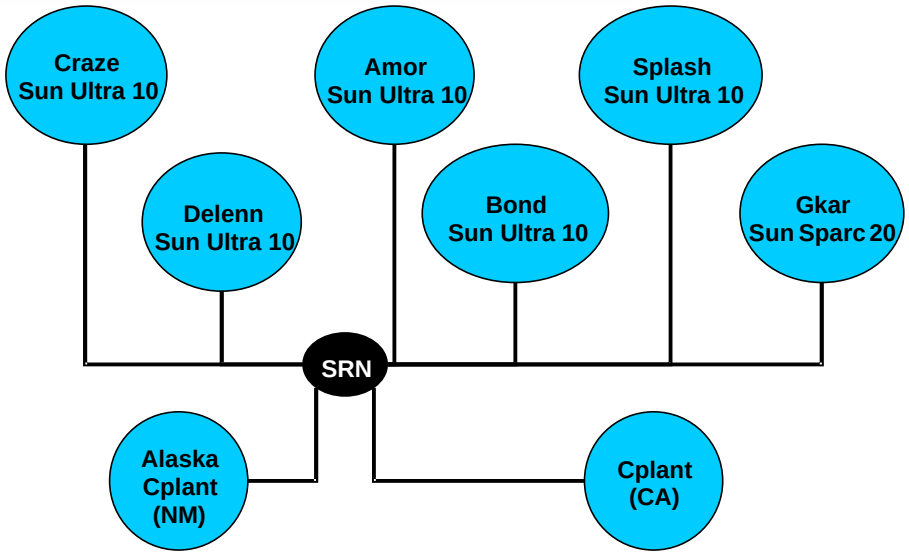

Figure 1. Schematic of DRM's Sandia Restricted Network(SRN) testbed.

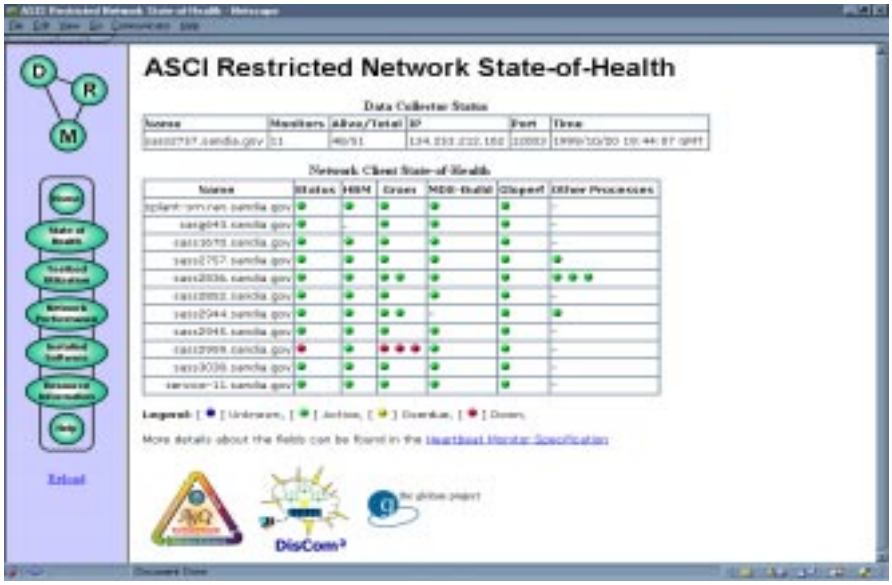

Figure 2. Web-based state-of-health monitoring. DRM clients can check the current status of grid resources.

CORBA Services

After DRM firmly established a Globus resource management foundation, DRM designed and implemented a Common Object Request Broker Architecture (CORBA) service layer on top of Globus. CORBA presents a standard way to componentize legacy or enterprise software. Because CORBA is platform and language independent, existing problem solving environments (PSEs) may easily integrate with DRM services. A product design environment and an optimization tool have been integrated with DRM services. Figure 3 illustrates a generic job submission PSE the DRM team implemented in Java.

The essential components of the DRM CORBA service layer are the broker and work manager. The broker uses the Globus Metacomputing Directory Service (MDS), a grid information base, to decide which resource is the best choice based on default matching algorithms or client-supplied matching constraints. DRM clients interface with the work manager. The work manager receives tasks from clients and queries the broker for a resource match. The work manager uses attributes of the selected resource to construct a concrete job specification that Globus runs. Figure 4 shows the interactions between DRM components.

Future Work

A primary FY00 goal for DRM is DOE security accreditation of the Globus and CORBA infrastructure. DRM will be deployed in the Tri-Lab complex and used for remote submission to the 10 TOps machine at LLNL. A coupled computation and visualization scenario will be used to develop and demonstrate new DRM services such as advance reservation, co-scheduling, and workflow management. More complex task specifications will be defined using eXtensible Markup Language (XML), which will simplify the CORBA service interface. It is expected that XML will play an important role for defining tasks in the Grid Community. The DRM team will continue to stay active in the Grid Forum. We also look forward to collaborating with many existing PSEs that want to take advantage of DRM's sophisticated resource management services.

For more information contact:

Sandia National Laboratories
Larry Ellis, Ljellis@sandia.gov

Lawrence Livermore National Laboratories
Morris Jette, Jette1@llnl.gov

Los Alamos National Laboratory
Jerry Melendez, Kjm@lanl.gov

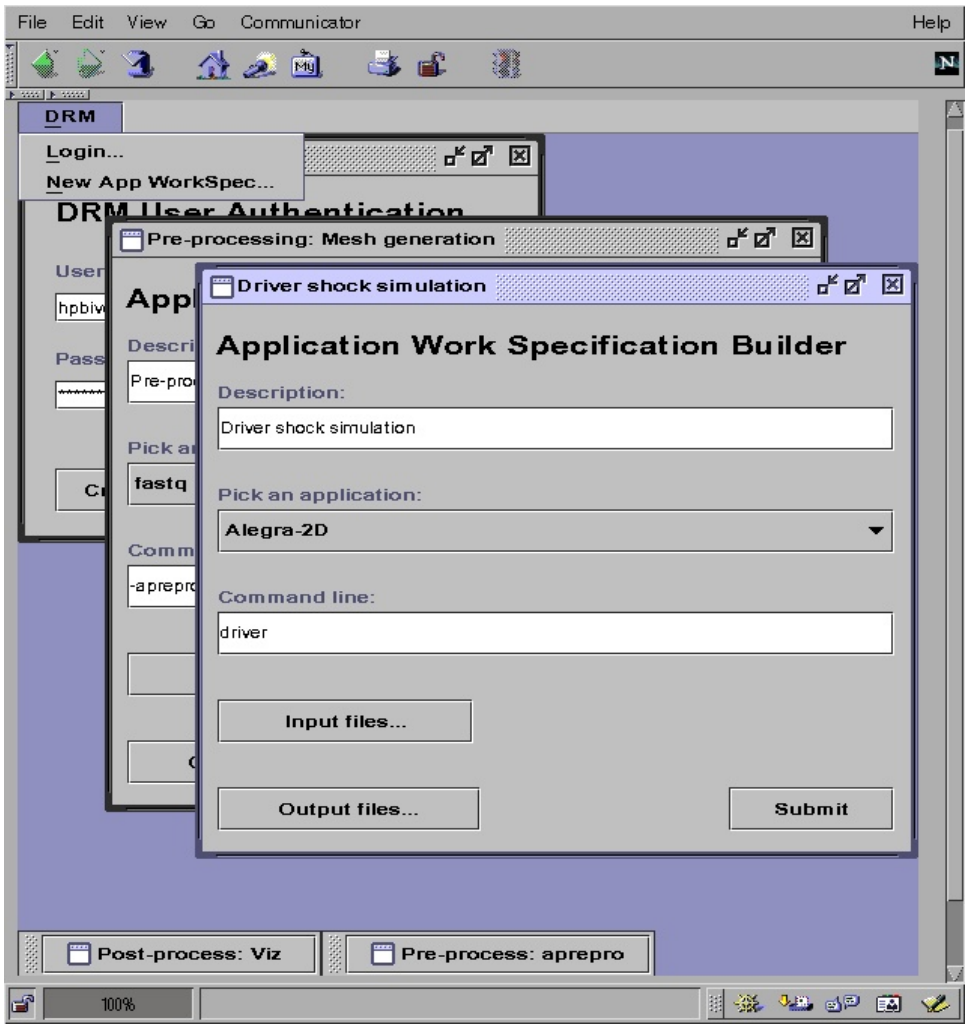

Figure 3. A simple job submission PSE using DRM CORBA Services. In this example, the user wishes to run a variety of different applications that perform pre- and post processing and numerical physics simulation. Note that users only need to specify an application name, commandline arguments, and input files. Because applications are a resource published in the MDS, the broker can find all the machines that support that software and pick the one with the lowest CPU load to run on. Here, the application Alegra-2D is a software resource that is well advertised in the MDS. Grid resources that support Alegra-2D will have a software entry for it. Software attributes indicate where the application physically resides on the resource, version, and any platform/machine specific information needed to run the code.

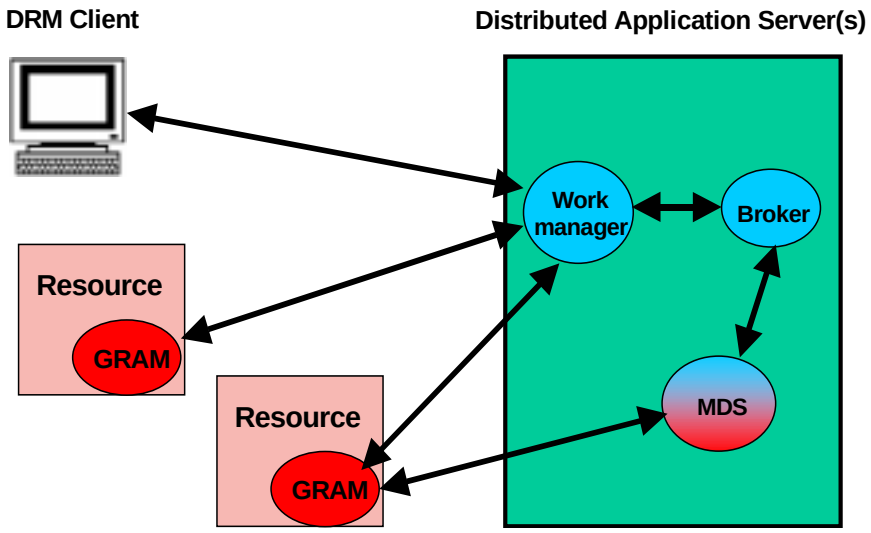

Figure 4. DRM's distributed object architecture. Globus Resource Allocation Managers (GRAMs) receive resource specifications from DRM work managers. When the GRAM allocates a resource, the state of the resource gets updated in the MDS.

ASCI, The Accelerated Strategic Computing Initiative Program, includes participants from Los Alamos National Laboratory, Lawrence Livermore National Laboratory, and Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.